arXiv:1605.07924v1 [stat.AP] 25 May 2016

# Modelling and Bayesian analysis of the Abakaliki Smallpox Data

Jessica E. Stockdale, Theodore Kypraios, Philip D. O'Neill

School of Mathematical Sciences, University of Nottingham

**Abstract**

The celebrated Abakaliki smallpox data have appeared numerous times in the epidemic modelling literature, but in almost all cases only a specific subset of the data is considered. There is one previous analysis of the full data set, but this relies on approximation methods to derive a likelihood. The data themselves continue to be of interest due to concerns about the possible re-emergence of smallpox as a bioterrorism weapon. We present the first full Bayesian analysis using data-augmentation Markov chain Monte Carlo methods which avoid the need for likelihood approximations. Results include estimates of basic model parameters as well as reproduction numbers and the likely path of infection. Model assessment is carried out using simulation-based methods.

## Keywords

Smallpox, Bayesian inference, Markov chain Monte Carlo, Stochastic epidemic model, Abakaliki

## 1 Introduction

In 1967, an outbreak of smallpox occurred in the Nigerian town of Abakaliki. The vast majority of cases were members of the Faith Tabernacle Church (FTC), a religious organisation whose members refused vaccination. A World Health Organization (WHO) report (Thompson and Foege , 1968) describes the outbreak, with information on not only the time series of case detections but also their place of dwelling (compound), vaccination status, and FTC membership. The outbreak has inherent historical interest as it occurred during the WHO smallpox eradication programme initiated in 1959. Although smallpox was declared eradicated in 1980, it regained attention as a potential bioterrorism weapon in the early 2000s (see e.g. Gani and Leach (2001), Meltzer *et al.* (2001) and Halloran *et al.* (2002))

and continues to be of interest due to concerns about its re-emergence or synthesis (see e.g. Henderson and Isao (2014), Eto *et al.* (2015), WHO (2015) and references therein). Public health planning for potential future smallpox outbreaks requires estimates of the parameters governing disease transmission, and thus being able to accurately obtain such quantities from available data is of considerable importance.

Within mathematical infectious disease modelling, the Abakaliki smallpox data set has been frequently cited, the first appearance being Bailey and Thomas (1971). The data are almost always used to illustrate new data analysis methodology, but in virtually all cases most aspects of the data are ignored apart from the population of 120 FTC individuals and the case detection times, while the models used are not particularly appropriate for smallpox (see for example Becker (1976), Yip (1989), O'Neill and Roberts (1999), O'Neill and Becker (2001), Huggins *et al.* (2004), Boys and Giles (2007), Lau and Yip (2008), Clancy and O'Neill (2008), Kypraios (2009), Shanmugan (2011), Xiang and Neal (2014), McKinley *et al.* (2014), Golightly *et al.* (2014), Oh (2014), Xu *et al.* (2016) and references therein). Ray and Marzouk (2008) use a more realistic smallpox model and take account of the compounds where individuals lived, but again ignore all non-FTC individuals.

The main objective of this paper is to present a Bayesian analysis of the full data set. To our knowledge, the only previous analysis of the full Abakaliki data is in Eichner and Dietz (2003), where the authors define a stochastic individual-based transmission model that considers not only the case detection times but also the other aspects of the data. Their model takes account of the population structure, the disease progression for smallpox, the vaccination status of individuals and the introduction of control measures during the outbreak. The model parameters are then estimated by constructing and maximising a likelihood function which is itself constructed using various approximations. Specifically, the true likelihood of the observed data given the model parameters is practically intractable, since it involves integrating over all possible unobserved events, such as the times at which individuals become infected. Eichner and Dietz tackle this problem by first using a back-calculation method to approximate the distribution of unobserved event times for a given individual, and then by making various assumptions about independence between individuals in order to construct an approximate likelihood function.

An alternative solution to the intractable likelihood problem is to use data-augmentation methods to produce an analytically tractable (and correct) likelihood, which can then be incorporated in a Bayesian estimation framework by using Markov chain Monte Carlo (MCMC) methods along the lines described in O'Neill and Roberts (1999) and Gibson and Renshaw (1998). We adopt this approach to carry out a full Bayesian analysis of the Abakaliki smallpox data, whilst also assessing how well the Eichner and Dietz approximation method works in this setting. Our approach provides results which can be directly compared with those of Eichner and Dietz, specifically estimates of model parameters, estimates of associated quantities of interest such as reproduction numbers, and the sensitivity of the results to the disease progression assumptions. In addition, we also estimate quan-

tities derived via data-augmentation, such as who-infected-whom and the time of infection for each individual, carry out various forms of model assessment to see how well the model fits the data, and explore particular aspects of the model via simulation. None of these additional elements feature in the Eichner and Dietz analysis.

The paper is structured as follows. In section 2 we describe the data, stochastic transmission model and method of inference. Section 3 contains results and details of sensitivity analysis and model-checking procedures. We finish with discussion in Section 4. The supplementary material contains details of some likelihood calculations and the MCMC algorithm.

# 2 Data, model and inference methods

The outbreak is described in detail in Thompson and Foege (1968) and Eichner and Dietz (2003). There were 32 cases in total, 30 of which were FTC members. All of the infected individuals lived in compounds, which were typically one-storey dwellings built around a central courtyard, and capable of housing several families. The FTC members frequently visited one another and were somewhat isolated from the rest of the community, which is one reason why most previous data analyses only consider FTC members. Although FTC members refused vaccination, many of them had been vaccinated prior to joining FTC as described below.

## 2.1 Abakaliki smallpox data

Table 1 contains details of the 32 cases of smallpox recorded during the outbreak, specifying the date of onset of rash, compound identifier, FTC membership status and vaccination status. Note that we set a timescale by setting day zero of the outbreak to be the first onset of rash date. The composition of the affected compounds is provided in Table 2, where the total numbers of vaccinated and non-vaccinated FTC and non-FTC members within each compound are listed. Note that on day 25, four FTC individuals from compound 1 (three vaccinated and one non-vaccinated) moved to compound 2. In addition, quarantine measures were put in place in Abakaliki, but not until part way through the outbreak. The exact time these measures were introduced was not recorded.

| Case number | Day of onset of rash | Compound | Confession | Vaccination |
|---|---|---|---|---|
| 0 | 0 | 1 | FTC | No |
| 1 | 13 | 1 | FTC | No |
| 2 | 20 | 1 | FTC | No |
| 3 | 22 | 1 | FTC | No |
| 4 | 25 | 1 | FTC | No |
| 5 | 25 | 1 | FTC | No |
| 6 | 25 | 1 | FTC | No |
| 7 | 26 | 2 | FTC | Yes |
| 8 | 30 | 2 | FTC | Yes |
| 9 | 35 | 1 | FTC | No |
| 10 | 28 | 4 | FTC | No |
| 11 | 40 | 5 | FTC | No |
| 12 | 40 | 1 | FTC | No |
| 13 | 42 | 1 | FTC | No |
| 14 | 42 | 1 | FTC | No |
| 15 | 47 | 1 | FTC | No |
| 16 | 50 | 5 | FTC | No |
| 17 | 51 | 2 | FTC | No |
| 18 | 55 | 1 | FTC | No |
| 19 | 55 | 2 | FTC | No |
| 20 | 56 | 6 | Non | Yes |
| 21 | 56 | 5 | FTC | Yes |
| 22 | 57 | 2 | FTC | Yes |
| 23 | 58 | 7 | FTC | No |
| 24 | 60 | 4 | FTC | No |
| 25 | 60 | 2 | FTC | No |
| 26 | 61 | 2 | FTC | No |
| 27 | 63 | 8 | Non | Yes |
| 28 | 66 | 3 | FTC | No |
| 29 | 66 | 9 | FTC | No |
| 30 | 71 | 5 | FTC | No |
| 31 | 76 | 2 | FTC | Yes |

Table 1: Smallpox cases in Abakaliki, Nigeria during 1967, taken from Thompson and Foege (1968). Compounds listed are those before the move of cases 7 and 8, and 2 other uninfected individuals, on day 25 from compound 1 to compound 2.

| Compound | FTC | | | Non-FTC | | |
|---|---|---|---|---|---|---|
| | Vaccinated | Nonvaccinated | $n_{c,FTC}$ | Vaccinated | Nonvaccinated | $n_{c,non}$ |
| 1 | 18 | 15 | 33 | 0 | 0 | 0 |
| 2 | 9 | 5 | 14 | 1 | 0 | 1 |
| 3 | 2 | 8 | 10 | 0 | 0 | 0 |
| 4 | 2 - $i_4$ | 2 + $i_4$ | 4 | 28 + $i_4$ | 1 - $i_4$ | 29 |
| 5 | 4 - $i_5$ | 3 + $i_5$ | 7 | 13 + $i_5$ | 2 - $i_5$ | 15 |
| 6 | 0 | 0 | 0 | 40 | 3 | 43 |
| 7 | 4 - $i_7$ | 1 + $i_7$ | 5 | 12 + $i_7$ | 3 - $i_7$ | 15 |
| 8 | 0 | 0 | 0 | 37 | 5 | 42 |
| 9 | 0 | 1 | 1 | 26 | 6 | 32 |
| Sum 1-9 | 35 | 39 | 74 | 161 | 16 | 177 |
| Outside | 46 x 35/74 | 46 x 39/74 | 46 | 30903 x 161/177 | 30903 x 16/177 | 30903 |
| Total | | | 120 | | | 31080 |

Table 2: Composition of the compounds affected by smallpox in Abakaliki, Nigeria during 1967, as described in Thompson and Foege (1968). Since we do not have complete vaccination status for all compounds, we use $i_4, i_5, i_7$ to allow for different possible configurations. The total number of vaccinated individuals is known and so $i_4 + i_5 + i_7 = 4$. It is also known that $i_4 \in \{0, 1\}, i_5 \in \{0, 1, 2\}$ and $i_7 \in \{1, 2, 3\}$. Note: this table displays the compound composition after the move of the 4 individuals from compound 1 to compound 2 on day 25.

## 2.2 Stochastic transmission model

We suppose that the residents of Abakaliki form a closed population with $N = 31,200$ individuals, labelled $0, 1, \ldots, N-1$. Individuals $0, 1, \ldots, n_{com} - 1$ are those inside the compounds, where $n_{com} = 251$ is the number of people within the compounds. Any individual $k = 0, \ldots, N-1$ may be categorised as type $(c_k, f_k)$, where (i) $c_k = 1, \ldots, 9$ is the compound of $k$, with $c_k = 0$ if $k$ is outside the compounds, and (ii) $f_k$ is $k$'s confession; FTC or non-FTC. These types may lead to differences in the mixing behaviour of individuals, but otherwise individuals are considered to be identical in their susceptibility to smallpox and their ability to infect others.

We now describe a stochastic disease-transmission model for the spread of smallpox throughout the population of Abakaliki. This model is essentially the same as that described in Eichner and Dietz (2003), and is a variant of an SEIR (Susceptible-Exposed-Infective-Removed) model. At any given time $t$ each individual in the population will be in any one of six states, namely susceptible, exposed, with fever, with rash, quarantined or removed. For $j = 0, \ldots, N-1$, let $e_j$, $i_j$, $r_j$, $q_j$, $\tau_j$ denote, respectively, the times of exposure, fever, rash, quarantine and recovery for individual $j$. Any susceptible individual may become exposed, as described below, at which point they enter an incubation (or latent) period. They next enter the fever stage of the disease, at which point they become infectious and may hence infect others. During the rash stage which follows, the individual remains infectious but with a potentially different level of infectivity. We define the infectious period to be the combined time spent in the fever and rash stages. Infectious individuals will either become removed (namely recovery or death; we do not distinguish these) or isolated, in which the individual is quarantined and henceforth unable to infect others. Control measures, in which cases are placed into isolation soon after detection, are introduced part way through the outbreak at time $t_q$. We do not allow re-infection, so that individuals who have been infected cannot become susceptible again. The epidemic continues until there are no infectious or exposed individuals remaining in the population, at which point each person will either still be susceptible, or will have been quarantined/removed.

The lengths of time spent in each stage of the disease for different individuals are assumed to be mutually independent random variables with specified distributions, the parameter values of which are assumed known. We adopt the assumptions of the Eichner and Dietz model, so that the incubation period, fever period and rash period all have Gamma distributions with values as shown in Table 3. If quarantine measures have been introduced, then an individual may be put into isolation after a random delay following their rash onset date. Specifically, we define the quarantine time of individual $j$ as $q_j = \max(r_j, t_q) + \Gamma(2, 2)$, where $\Gamma(\mu, \sigma)$ denotes a gamma distributed random variable with mean $\mu$ and standard deviation $\sigma$. This means that no quarantining occurs prior to time $t_q$, after which it takes an average of two days for a detected individual to be placed in isolation. Note that both removal and quarantining of an individual are equivalent in the sense that both mean the individual can

| | Mean (days) | Standard deviation (days) |
|---|---|---|
| Period before fever | $\mu_I = 11.6$ | $\sigma_I = 1.9$ |
| From fever to rash | $\mu_F = 2.49$ | $\sigma_F = 0.88$ |
| From rash until recovery | $\mu_R = 16.0$ | $\sigma_R = 2.83$ |
| From rash to quarantine or from $t_q$ to quarantine | $\mu_Q = 2.0$ | $\sigma_Q = 2.00$ |

Table 3: Durations of disease stages in the smallpox model. The time until quarantine changes after the introduction of control measures as described in the text.

no longer infect others, but we include both in the model for clarity, and also for comparison with the Eichner and Dietz model. Note also that an infected individual will have both a removal time and a quarantine time, and both quantities appear in the likelihood function as explained later.

We assume that the epidemic is initiated by a single exposed individual, whom we label $\kappa$. The epidemic thus begins at time $e_\kappa$ with the exposure of the initial infective $\kappa$. Recall that the infectious period is defined in two parts: the fever period and the rash period, during each of which an individual will be infectious, but at potentially different rates. During their rash period, an individual $j$ will have contacts with other members of their compound who are of the same confession at times given by the points of a Poisson process of rate $\lambda_h$ per day. Individuals outside of the nine compounds do not have such contacts. In addition, FTC individuals will have contacts at rate $\lambda_f$ per day with other FTC individuals and contacts at rate $\lambda_a$ per day with anybody (including FTC individuals) in the population. Non-FTC individuals are assumed to have contacts with anybody in the population at rate $\lambda_a + \lambda_f$ per day. This assumption is made to ensure that all individuals have the same average number of contacts per day outside of their own compound. During the fever period, contacts occur in exactly the same manner except that all rates are multiplied by a factor $b$, to account for a potential difference in infectivity during the fever period. In each case, the individual actually contacted is chosen uniformly at random from the pool of potential individuals in question. For example, contacts made by an individual with the entire population are drawn from the $N - 1$ other individuals. Note that this means that the individual-to-individual contact rate for such contacts is $\lambda_a/(N-1)$. Any contact from an infective to a susceptible results in immediate exposure of the susceptible. All of the Poisson processes describing contacts are assumed to be mutually independent.

In addition, a proportion of the population is vaccinated. The vaccination status of all but a few individuals within the compounds is known, and the proportions of FTC/non-FTC vaccinated individuals outside the compounds is assumed to be the same as inside.

However, vaccination is not necessarily effective: each recipient of the vaccine is completely protected with probability $v$, or remains completely susceptible with probability $1 - v$. Although the total number of vaccinated individuals is known, we do not have complete information on the composition of individuals with respect to vaccination status and FTC membership (see Table 2). There are five potential configurations of twelve individuals with unknown details to consider. For each individual in the population we will have a vaccination status, which is assumed known for most individuals, and a protection status, which is unknown.

### 2.3 Reproduction numbers

Within mathematical epidemic modelling, the so-called basic reproduction number is of primary importance. It can be defined as the average number of cases caused by a typical index case in a large population, and its value typically governs certain aspects of the epidemic such as the final number of cases or the probability of an epidemic dying out rapidly. From a mathematical viewpoint, reproduction numbers for stochastic epidemics are typically defined by allowing the population size to become infinite in an appropriate sense. For models featuring structured populations with different kinds of contact rates, reproduction numbers can be defined in various ways. Eichner and Dietz consider two such reproduction numbers, namely $R_0 = (\mu_R + b\mu_F)(\lambda_a + \lambda_f + \lambda_h)$ and $R_F = b\mu_F(\lambda_a + \lambda_f + \lambda_h)$. Here, $R_0$ is a reproduction number for an infected FTC member within the compounds, and can be interpreted as the average number of new infections such an individual would cause, under the assumption that the FTC population, compound populations and entire population are all large. Similarly, $R_F$ describes the average number of new infections caused by contacts made during the fever period of an FTC individual within the compounds. As discussed later, one drawback with such definitions for the Abakaliki data is that the compounds themselves are not particularly large. However, for comparison purposes we shall also consider $R_0$ and $R_F$ in our analysis.

### 2.4 Bayesian inference and MCMC

Our aim is to perform Bayesian inference for the unknown model parameters given the data, which consist of rash times for all infectives, knowledge of the population structure and vaccination status of individuals. We will use an MCMC algorithm to obtain approximate samples from the posterior density of the model parameters, namely the infection rates $\lambda_a$, $\lambda_f$ and $\lambda_h$, the vaccine efficacy $v$, the infectivity factor $b$ and the time quarantine measures were introduced, $t_q$. Our approach involves data augmentation, specifically involving the exposure, fever, removal and quarantine times of each infected individual, and also the protection statuses and unknown vaccination statuses. First we derive an expression for the

likelihood of the observed and augmented data. Let

$$\boldsymbol{\Phi} = (\kappa, e_\kappa, t_q, b, v, \lambda_a, \lambda_f, \lambda_h, \boldsymbol{\theta}, \mathbf{s})$$

where

$$\boldsymbol{\theta} = (\mu_I, \sigma_I, \mu_F, \sigma_F, \mu_R, \sigma_R, \mu_Q, \sigma_Q),$$

so that the components of $\boldsymbol{\theta}$ are parameters that are assumed known, and where $\mathbf{s} = (s_0, s_1, ..., s_{N-1})$, where for $i = 0, 1, \ldots, N-1$, $s_i$ is equal to 1 if individual $i$ is vaccinated and 0 if not. These vaccination statuses are assumed known for the majority of individuals, with a small number of exceptions, as shown in Table 2.

We define $\mathbf{e}$, $\mathbf{i}$, $\mathbf{q}$ and $\boldsymbol{\tau}$ as the unknown sets of exposure (not including the initial exposure $e_\kappa$), infection, quarantine and removal times, respectively. Similarly, we define $\mathbf{r}$ as the known set of rash times for all infectives. Then we define the augmented data as

$$\boldsymbol{\gamma} = (\mathbf{e}, \mathbf{i}, \mathbf{q}, \boldsymbol{\tau}, \mathbf{p}),$$

where $\mathbf{p} = (p_0, p_1, ..., p_{N-1})$ contains the unknown protection status of each individual, specifically with $p_i = 1$ if individual $i$ is protected, and $p_i = 0$ if they are not.

For an individual $j = 0, ..., N-1$ who is susceptible at time $t$ we define $\Lambda_j(t)$ as the infectious pressure acting upon them at time $t$, so that

$$\mathbb{P}(j \text{ is exposed in } (t, t+\delta t]) = \Lambda_j(t)\delta t + o(\delta t),$$

whilst for an individual $j$ who is no longer susceptible at time $t$ we set $\Lambda_j(t) = 0$. From the model definition, if $j$ is susceptible at time $t$ then $\Lambda_j(t)$ can be written as

$$\Lambda_j(t) = \sum_{k \in \mathcal{N}_{inf}(t)} m(k,t) \times \begin{cases} \frac{\lambda_a}{N-1} + \frac{\lambda_f \delta_f(k,j)}{n-1} + \frac{\lambda_h \delta_c(k,j;t)}{n_{c,f_k}(t)-1} & \text{if } f_k = \text{FTC}, \\ \frac{\lambda_a+\lambda_f}{N-1} + \frac{\lambda_h \delta_c(k,j;t)}{n_{c,f_k}(t)-1} & \text{otherwise,} \end{cases}$$

where

$$m(k,t) = \begin{cases} b & \text{if } i_k \le t < r_k, \\ 1 & \text{if } r_k \le t < \min(\tau_k, q_k), \\ 0 & \text{otherwise,} \end{cases}$$

and (i) $\delta_f(j,k) = 1$ if both $k$ and $j$ are FTC and zero otherwise; (ii) $\delta_c(j,k;t) = 1$ if both $k$ and $j$ live in the same compound at time $t$ and are of the same confession, and zero otherwise; (iii) $n = 120$, the number of FTC individuals in the population; (iv) $n_{c,f_j}(t)$ is the number of individuals in $j$'s compound of the same confession as $j$ at time $t$, including $j$ themselves, and (v) $\mathcal{N}_{inf}(t)$ is the set of individuals infective at time $t$.

We denote the likelihood of the data $\mathbf{r}$ given the model parameters $\boldsymbol{\Phi}$ as $\pi(\mathbf{r}|\boldsymbol{\Phi})$. This is practically intractable since its evaluation involves integrating over all possible unobserved

events. We instead proceed by augmenting the data $\mathbf{r}$ with $\boldsymbol{\gamma}$ to obtain the tractable likelihood

$$\begin{aligned}\pi(\mathbf{r},\boldsymbol{\gamma}|\boldsymbol{\Phi}) &= \left(\prod_{j\in\mathcal{N}_{inf}}\Lambda_j(e_j-)\right)\times \mathrm{e}^{-\int_{e_\kappa}^{T}\Lambda(t)dt}\\ &\quad\times\prod_{j\in\mathcal{N}_{inf}} f_I(i_j-e_j)f_F(r_j-i_j)f_R(\tau_j-r_j)f_Q(q_j-\max(r_j,t_q))\\ &\quad\times\upsilon^{\sum_{r=0}^{N-1}p_r s_r}(1-\upsilon)^{\sum_{r=0}^{N-1}(1-p_r)s_r},\end{aligned} \tag{1}$$

where (i) for $t\geq e_\kappa$,

$$\Lambda(t)=\sum_{j=0}^{N-1}\Lambda_j(t)$$

denotes the total pressure acting on all susceptible individuals at time $t$; (ii) $\Lambda_j(e_j-)=\lim_{t\uparrow e_j}\Lambda_j(t)$ is the pressure on $j$ just before their exposure; (iii) $\mathcal{N}_{inf}$ is the set of individuals who ever become infected; (iv) $T$ is the end of the epidemic, i.e. the first time at which no infectives or exposed individuals remain in the population (in practice we set $T$ equal to the final rash time); and (v) $f_A$, for $A=(I,\ F,\ R,\ Q)$, is the probability density function of a $\Gamma(\mu_A,\sigma_A)$ distribution. The augmented likelihood function in (1) is of a fairly standard form (see e.g. O'Neill and Becker (2001)) and contains the following components. The first product term accounts for the exposure of each of the observed cases and the exponential term gives the probability of individuals avoiding infection (either until they become infected, or throughout the entire epidemic). The second product term gives the likelihood of the times spent in each of the disease progression states for each individual who ever becomes infected. The final terms give the probability of the protection statuses for all individuals in the population.

One practical drawback with our data augmentation scheme as it stands is that it includes protection statuses $p_i$ for all $N=31,200$ individuals in the population. However, it is possible to integrate out these parameters for all individuals outside the compounds, essentially because the number of protected individuals follows a Binomial distribution. The calculations are fairly lengthy and so are provided in the supplementary material.

By Bayes' Theorem, the posterior density of interest is

$$\pi(\boldsymbol{\Phi},\boldsymbol{\gamma}|\mathbf{r})\propto\pi(\mathbf{r},\boldsymbol{\gamma}|\boldsymbol{\Phi})\pi(\boldsymbol{\Phi}),$$

where $\pi(\boldsymbol{\Phi})$ denotes the prior density of $\boldsymbol{\Phi}$. We assume *a priori* independence of the components of $\boldsymbol{\Phi}$ so that

$$\pi(\boldsymbol{\Phi})=\pi(\kappa)\pi(t_q)\pi(b)\pi(\upsilon)\pi(\lambda_a)\pi(\lambda_f)\pi(\lambda_h)\pi(\boldsymbol{\theta})\pi(\mathbf{s}).$$

We set $\lambda_a$, $\lambda_f$ and $\lambda_h$ to have $\Gamma(10^3, 10^6)$ prior distributions which corresponds to vague prior assumptions for these parameters, set $v$, $b$ and $t_q$ to have uniform prior distributions on $(0, 1)$, $(0, \infty)$ and $(0, \infty)$ respectively, and set $\kappa$ to have a discrete uniform distribution over all infected individuals. Since $\boldsymbol{\theta}$ is assumed known, $\pi(\boldsymbol{\theta})$ is just a point mass. Finally, $\pi(\mathbf{s})$ consists of a point mass at the known vaccination statuses with a uniform distribution over the five possible configurations of twelve unknown vaccination statuses as shown in Table 2.

We use an MCMC algorithm to produce samples of the parameters of interest from the target posterior distribution, updating both the model parameters and the unknown event times as well as protection statuses and vaccination combinations. The algorithm is non-standard, and although it is similar in principle to that in O'Neill and Roberts (1999), in practice it is far more complex and involves considerable book-keeping. Full details of the algorithm are provided in the supplementary material.

# 3 Results and analysis

## 3.1 Parameter estimates and reproduction numbers

Table 4 contains posterior summaries for the model parameters along with the corresponding maximum likelihood estimates from Eichner and Dietz' approximate likelihood method, and Figure 1 contains corresponding density plots, scatter plots and posterior correlation coefficients. Our estimates are fairly similar to those of Eichner and Dietz, and in particular the posterior modal values for the six basic model parameters are quite close. Although they represent different quantities, our posterior credible intervals and the confidence intervals of Eichner and Dietz are also quite similar. Our mean estimate of the basic reproduction number $R_0$ is 7.96, which is slightly higher than Eichner and Dietz' estimate of 6.87. Similarly, our estimate of the reproduction number for the fever period $R_F$ is 0.53 compared to Eichner and Dietz's 0.164. In this case the difference can be explained by the highly skewed posterior density for $b$, so that the mean and mode are clearly different. The scatter plots and correlation coefficients suggest that the basic model parameters can be separately estimated from the data and that the model is not over-parameterised.

## 3.2 Who infected whom

Since the MCMC algorithm involves imputation of all event times, it is straightforward to obtain estimates of the path of infection, i.e. who infected whom. Specifically, if an individual $j$ is subject to infectious pressure $\Lambda_j(t) = \sum_{k=1}^{m} a_k(t)$ at the time of their exposure, where $a_k(t)$ is the pressure from the $k$th of $m$ infectives at time $t$, then the probability that individual $k$ actually infected $j$ is simply $a_k(t)/\Lambda_j(t)$. Figure 2 shows the most likely infector for each observed case and also a greyscale plot which illustrates the associated

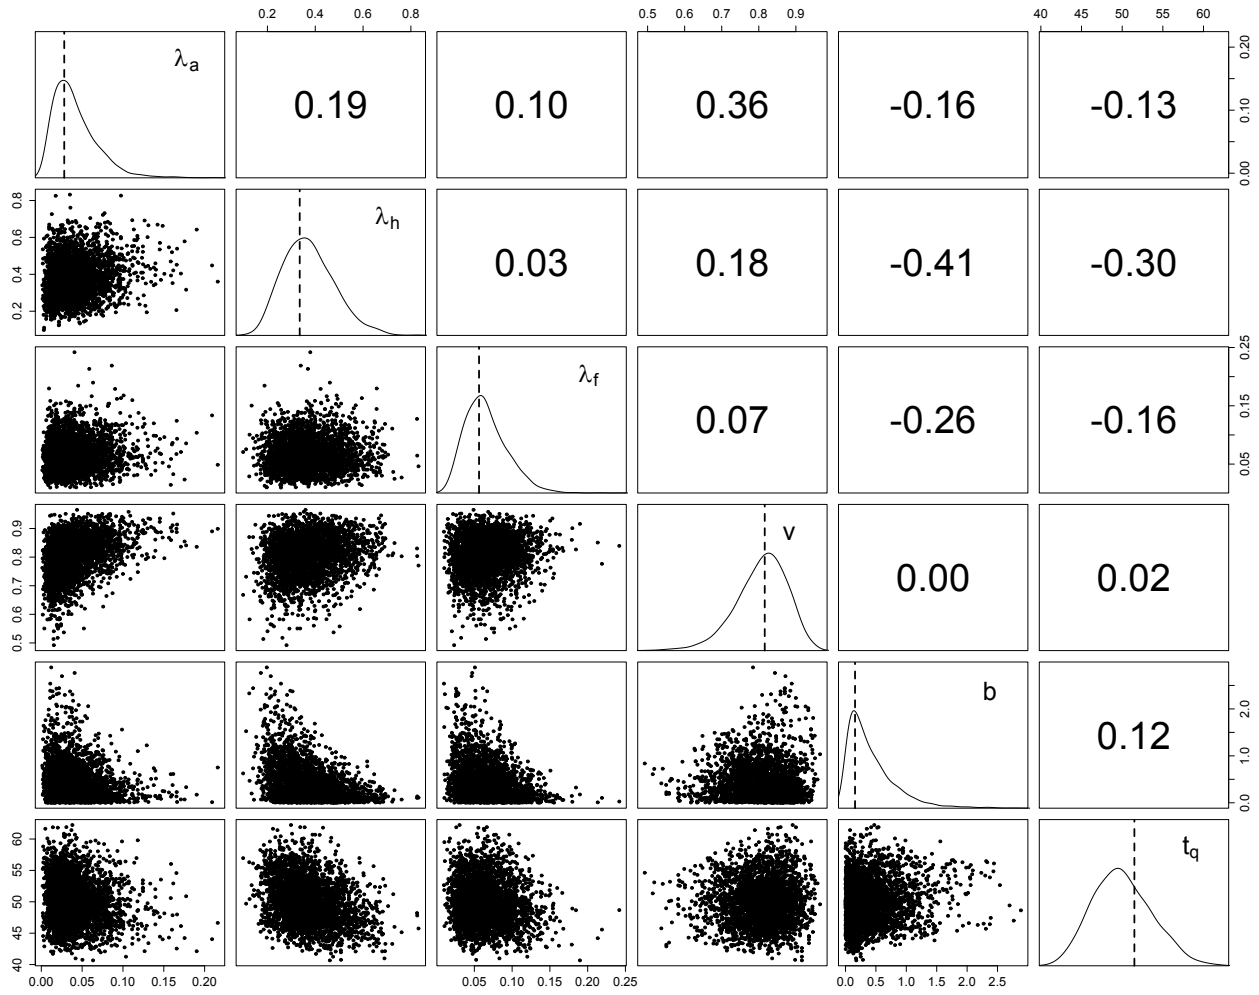


Figure 1: The diagonal shows posterior densities for the 6 parameters of interest, based on 100,000 MCMC samples. Dotted lines give the Eichner and Dietz estimates. The lower panels show scatter plots of the MCMC samples and the upper panels show posterior correlation coefficients.

| Parameter | Posterior mean | Posterior median | Credible interval | ED estimate | ED confidence interval |
|---|---|---|---|---|---|
| $\lambda_a$ | 0.041 | 0.035 | (0, 0.093) | 0.0281 | (0.00447, 0.101) |
| $\lambda_f$ | 0.063 | 0.059 | (0.009, 0.010) | 0.0562 | (0.0187, 0.127) |
| $\lambda_h$ | 0.358 | 0.349 | (0.150, 0.565) | 0.335 | (0.192, 0.527) |
| $v$ | 0.808 | 0.817 | (0.668, 0.947) | 0.816 | (0.644, 0.922) |
| $b$ | 0.522 | 0.374 | (0, 1.500) | 0.157 | (0, 1.89) |
| $t_q$ | 50.4 | 50.2 | (42.4, 58.4) | 51.5 | (44.7, 59.6) |
| $R_0$ | 7.96 | 7.79 | (4.33, 11.56) | 6.87 | (4.52, 10.1) |
| $R_F$ | 0.531 | 0.431 | (0, 1.364) | 0.164 | (0, 1.31) |

Table 4: Parameter estimates and equal-tailed 95% posterior credible intervals for the Abakaliki smallpox outbreak from the true likelihood approach, with results from Eichner and Dietz (2003) for comparison (ED). 100,000 MCMC samples were used. Here $R_0 = (\mu_R + b\mu_F)(\lambda_a + \lambda_f + \lambda_h)$ and $R_F = b\mu_F(\lambda_a + \lambda_f + \lambda_h)$.

uncertainty. We see that most infections occurred within compounds; note that individuals 7 and 8 moved from compound 1 to compound 2 during the outbreak and so in reality most of the infections caused by individual 8 were probably also within-compound. Most individuals give rise to one secondary case, but individuals 0 and 8 both cause multiple secondary cases. The greyscale plot shows that there is a modest degree of uncertainty around the identity of each infector.

### 3.3 Exposure times

Figure 3 illustrates the posterior distribution of the initial exposure time for each of the 32 cases. Generally speaking there is relatively little uncertainty, and most of the exposure times follow the ordering of the rash times, both features that are likely to be consequences of the distributions assumed for disease progression. The figure also shows the temporal aspects of the outbreak in terms of generations of infection: the first two generations (i.e. those infected by the index case, and those they infect) are clearly discernible, while the third and fourth generations are less distinct from each other, although (according to Figure 2) the fourth generation only contains two individuals. We see some groups of individuals with very similar exposure times and who are all infected by the same person, according to Figure 2, individuals 4, 5, 6 and 10, 11, 12 being two examples. Such clustering, more akin to a point-source outbreak, illustrates the high transmission potential for smallpox in close-contact settings. Finally, we comment that Eichner and Dietz (2003) also provide a plot showing likely exposure times, along with other event times, but that this is based entirely on back-calculation using the assumed disease progression model. In particular, their plot takes no account of the transmission model itself.

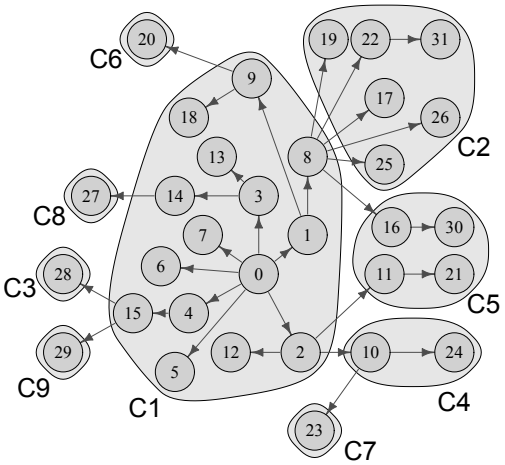


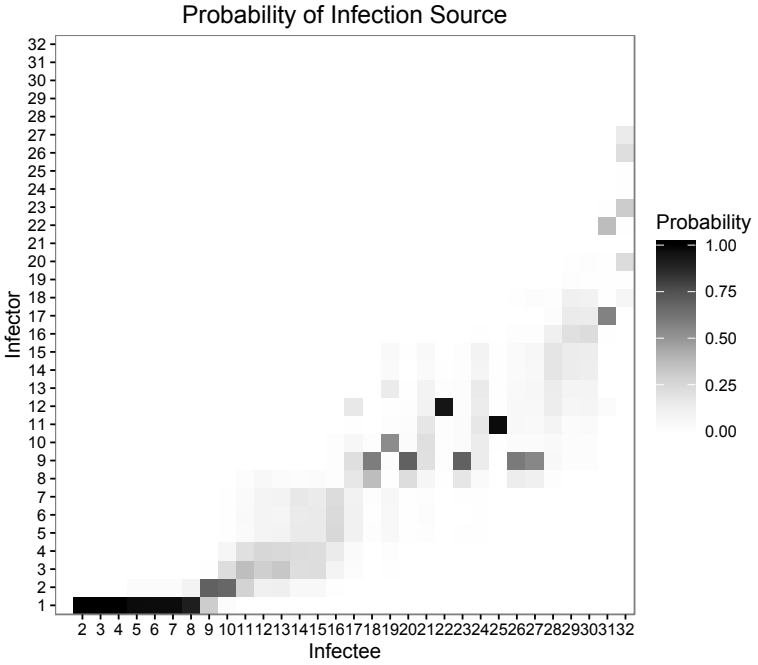


Figure 2: Estimation for who-infected-whom during the Abakaliki outbreak. The first figure shows the most likely infector for each individual, i.e. the infector shown has the highest posterior probability among all possible infectors. Labels C1, ..., C9 denote compounds 1, ..., 9, respectively. Individuals 7 and 8 moved into compound 2 during the outbreak; the figure shows the initial configuration. The second figure shows the posterior distribution of possible infectors for each infectee.

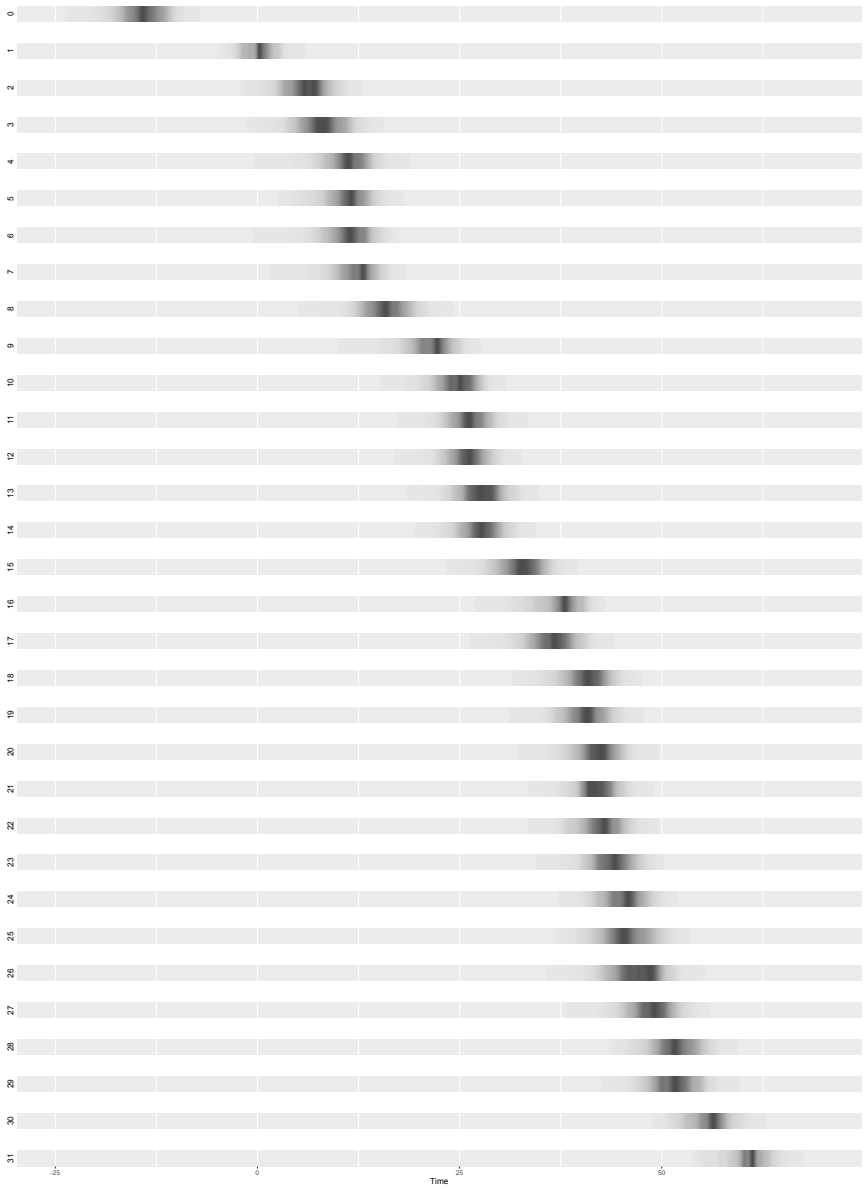


Figure 3: Heatmap showing the posterior distribution of time of exposure for each of the 32 cases in the Abakaliki outbreak.

### 3.4 Sensitivity analysis

We now briefly explore the sensitivity of our results to the underlying model assumptions, and in particular the assumed values for the periods of time spent in each disease state. Figure 4 displays posterior densities for the parameters of interest over a range of values taken from Meltzer *et al.* (2001) and Gani and Leach (2001). As might be expected, when $\mu_R$, the mean length of the rash period before removal, is reduced to make shorter average infectious periods, the estimates of the infection rate parameters increase to compensate. It is of interest to note that estimation of $R_0$ is somewhat sensitive to the choice of $\mu_R$. This is likely to be an artefact of the fact that that there are relatively few cases, the population structure and the introduction of control measures, since in a large uninterrupted outbreak we would expect $R_0$ to be more or less determined by the outbreak size. Figure 5 shows the effect of varying $\mu_Q$ and $\sigma_Q$, the parameters which govern the time taken for a case to be quarantined. Specifically, we consider the effect of halving or doubling the mean time, while keeping the coefficient of variation fixed at unity. Here we see relatively little impact, which is reassuring since we have very little information on which to base our modelling assumptions, in contrast to those which depend on more generic features of smallpox.

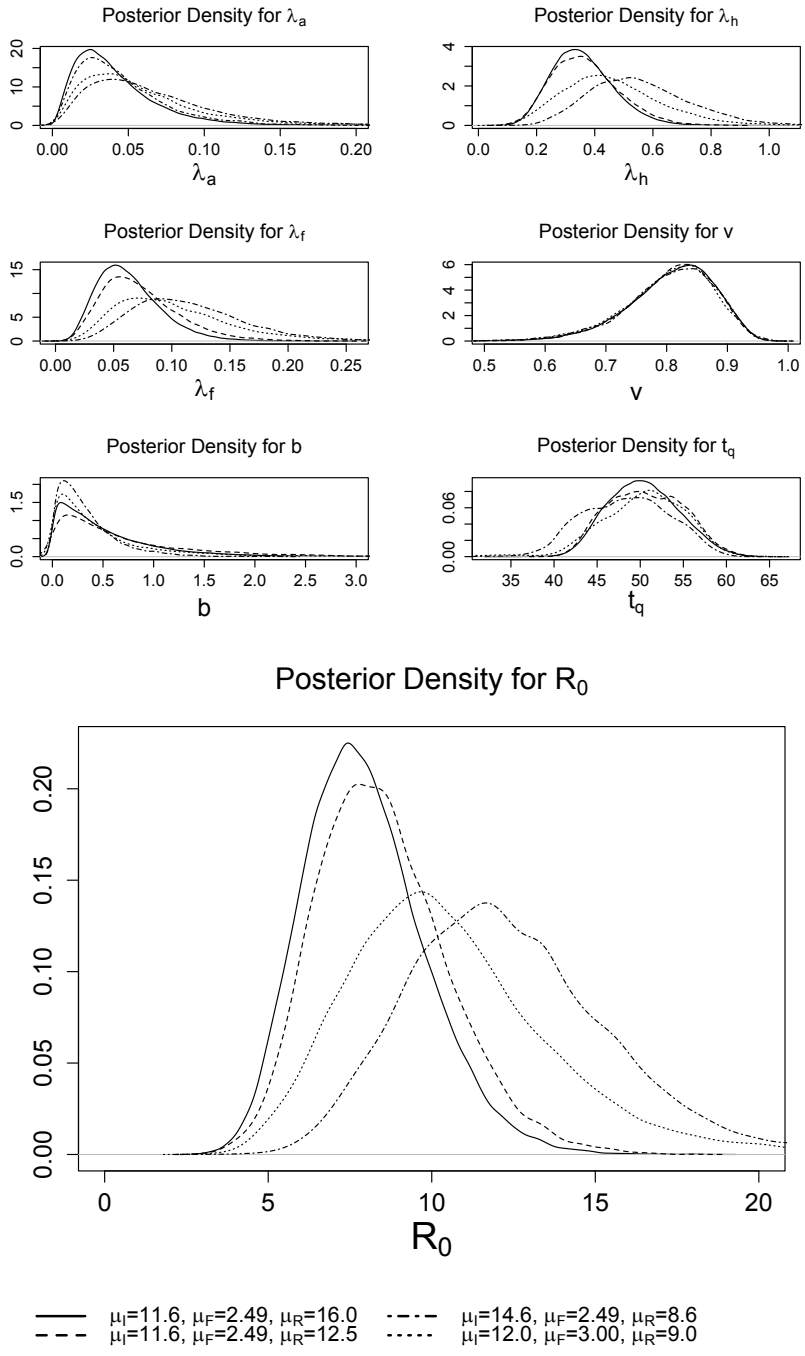


Figure 4: Posterior densities of the model parameters and $R_0$ using different mean durations of the disease stages. The solid line represents the baseline case as used in Eichner and Dietz and our primary analysis.

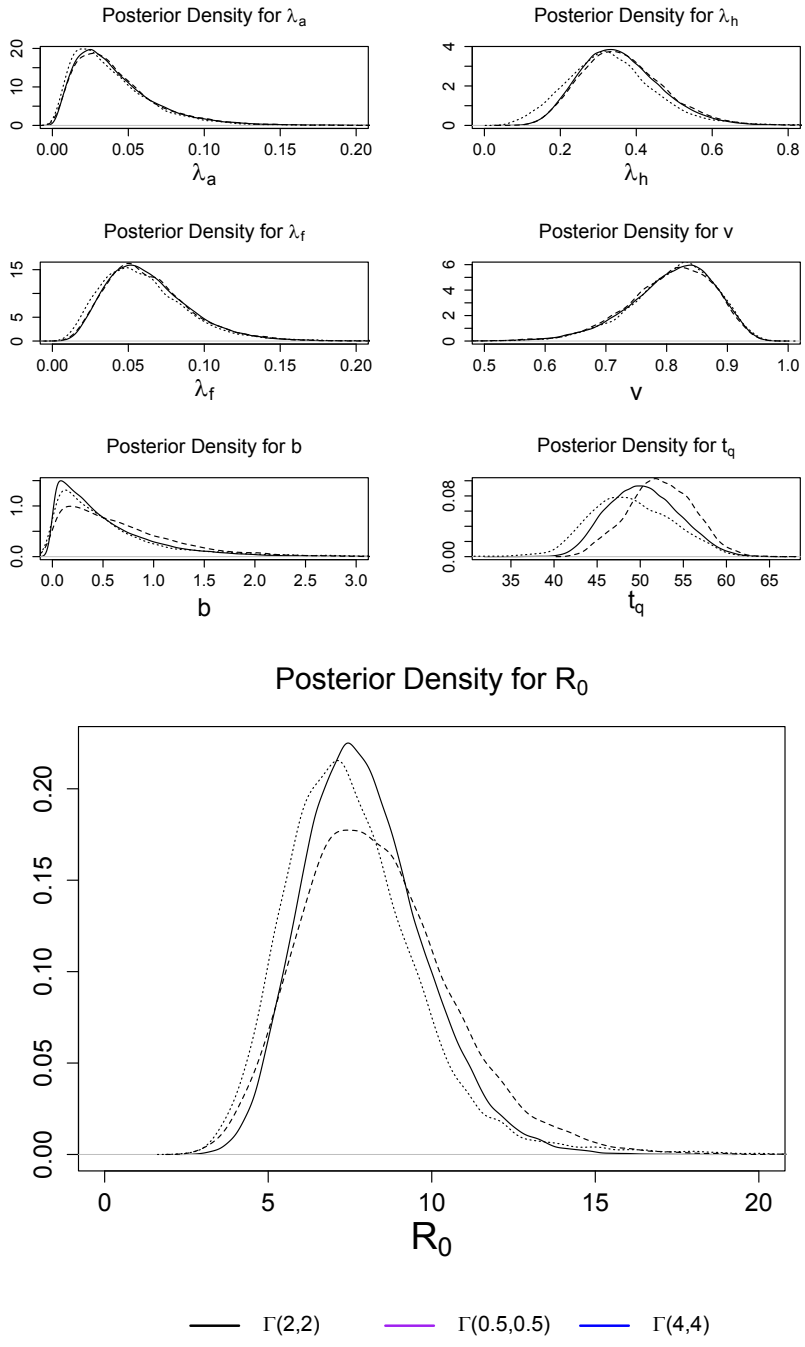


Figure 5: Posterior density of the model parameters and $R_0$ as the time taken to quarantine a detected case is varied.

### 3.5 Model assessment

In order to assess how well our model fits the data we use its posterior predictive distribution. Specifically, we take samples of the basic model parameters from the MCMC output and simulate the model forwards in time for each set of parameter values. We then compare various aspects of the observed data to the ensemble of simulations to see if the former aligns in some sense with the latter.

We start with the final size of the epidemic, i.e. the total number of cases. Figure 6 shows that although the observed final size (32) is not untypical of those produced by the model, it is some way from the mean (23.5) and mode of the distribution. This underestimation appears to be largely due to the fact that in the actual outbreak, four individuals, of whom two were infected, moved from compound 1 to compound 2, leading to new cases in compound 2. To account for this, Figure 6 also shows a histogram of the final size distribution among those simulated epidemics in which at least one of the four moving individuals was infected. It can be seen that this adjustment gives a better fit to the observed final size.

We next consider epidemic duration, defined as the length of time between the first case detection (rash) time and the last. Figure 7 shows a histogram of the durations of 5000 simulated outbreaks. The mean duration is 76.8 days, which is very similar to the Abakaliki outbreak (76 days). Including only those outbreaks in which infected individuals moved compound only increases the mean by a few days.

We next compare the simulated cumulative number of cases with the Abakaliki data. This is complicated by the fact that different simulated outbreaks usually have different total numbers of cases, and so to facilitate the comparison we consider only simulated outbreaks that have the same total number of cases (32) as the data. Figure 8 shows the results of 1000 simulations, from which it appears that the observed data are reasonably well captured by the model behaviour. To quantify this more precisely, we calculated a posterior predictive p-value as follows. Recall the chi-squared discrepancy measure (see e.g. Gelman *et al.* (1996)), which here takes the form

$$D(\mathbf{r}, \mathbf{\Phi}) = \sum_j \frac{\big(r_j - \mathbb{E}(r_j|\mathbf{\Phi})\big)^2}{\mathrm{Var}(r_j|\mathbf{\Phi})},$$

where $\mathbf{\Phi}$ denotes the model parameters and $\mathbf{r} = (r_1, \ldots, r_{32})$ denotes a vector of case-detection (rash) times. Note that neither the mean nor variance term are available analytically, and so in practice they are obtained via simulation: given $\mathbf{\Phi}$, we simulate epidemics repeatedly until we have a sample of size $M_1$, all with 32 cases. The mean and variance of the $j$th rash time is then estimated directly from this sample. Suppose now that we have $M$ samples from the posterior density of the model parameters, labelled $\mathbf{\Phi}^{(1)}, \ldots, \mathbf{\Phi}^{(M)}$. We use the $i$th sample to obtain a simulated epidemic with 32 cases and rash times $\mathbf{r}_i^{\mathrm{rep}}$

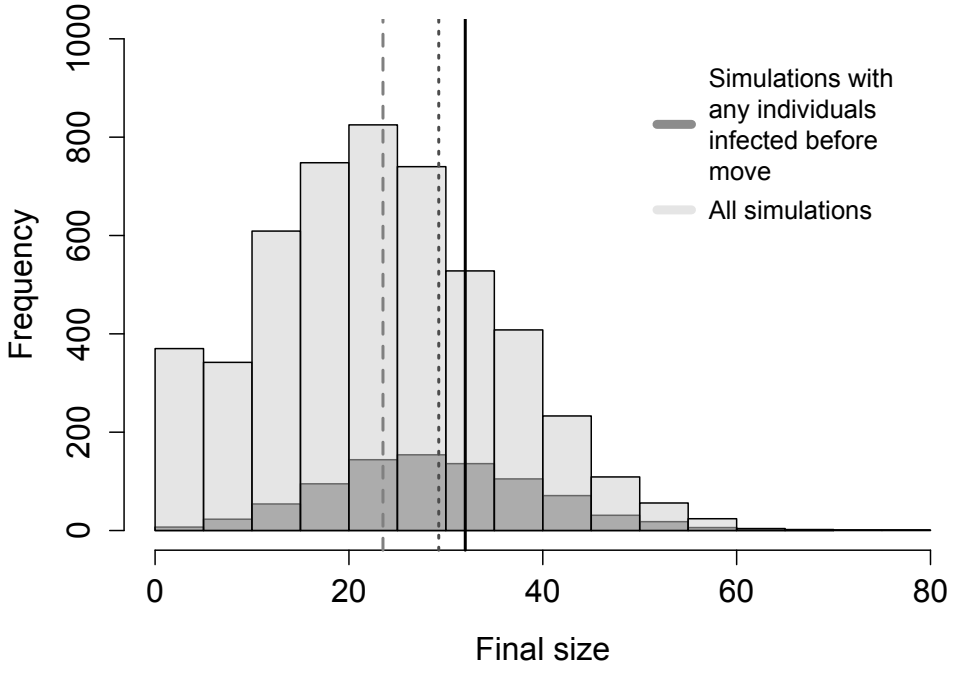


Figure 6: Final size of 5000 outbreaks simulated from posterior estimates (mean 23.5, dashed line), and the subset of 848 outbreaks in which at least one of the four individuals who moved compound was infected prior to the move (mean 29.3, dotted line). The solid line shows the observed final size (32).

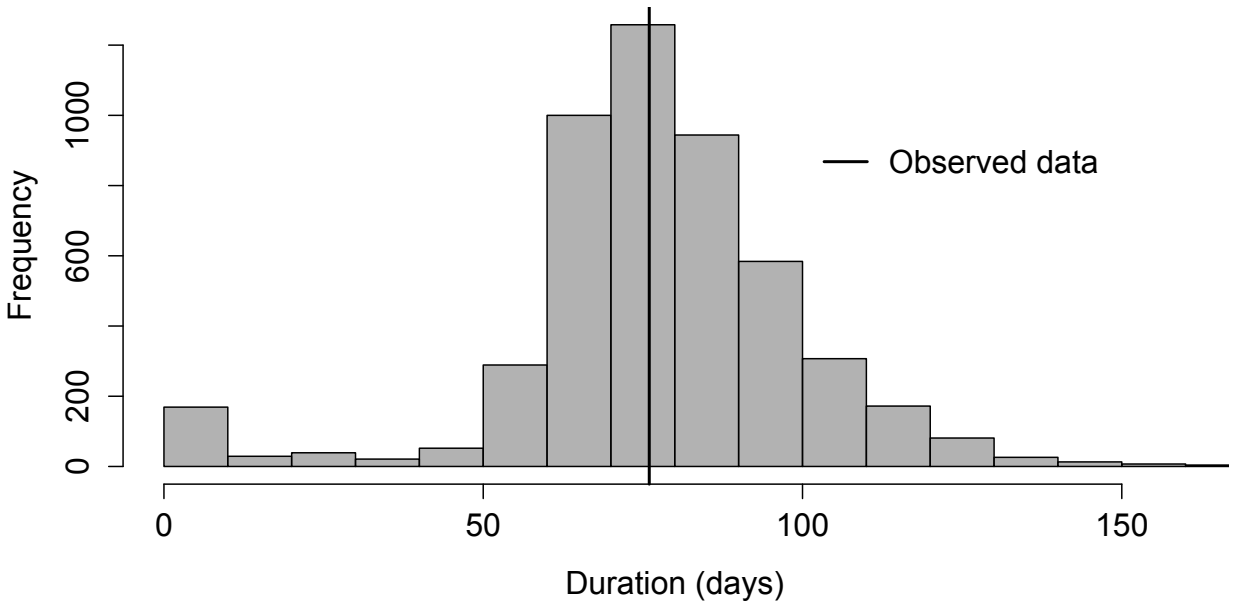


Figure 7: Duration of 5000 outbreaks simulated from posterior estimates, with the solid line showing the duration of the Abakaliki outbreak.

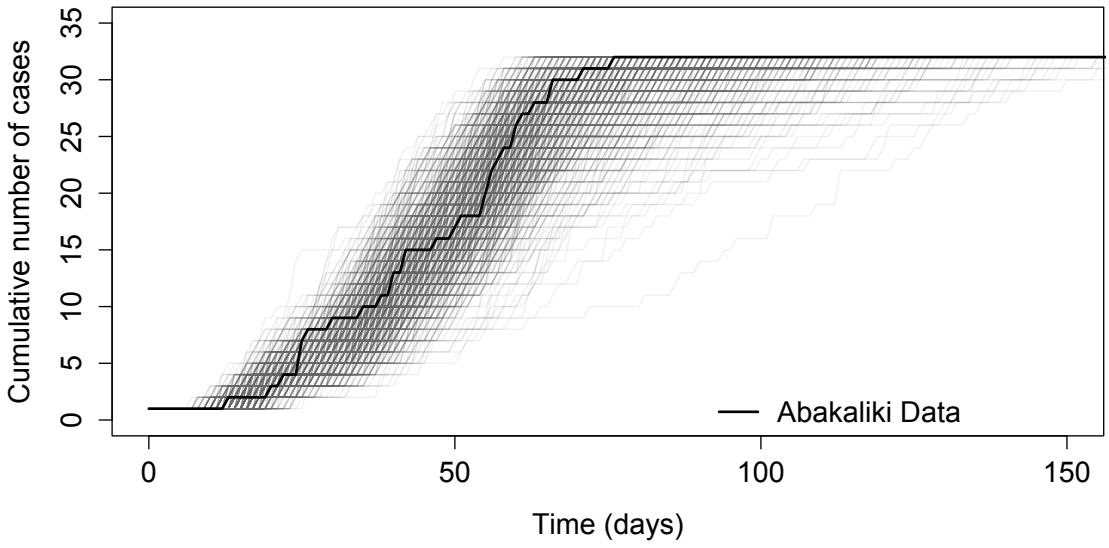


Figure 8: 1000 simulated epidemics each with a total of 32 cases. The solid line is the observed cumulative number of cases.

by repeatedly simulating epidemics until we obtain one with 32 cases. Denoting a typical simulation replicate $\mathbf{r}^{\text{rep}}$ and letting $\mathbf{r}^{\text{obs}}$ denote the observed rash times, the posterior predictive p-value is defined as

$$\begin{aligned} \text{ppp-value} &= \mathbb{P}\Big(D(\mathbf{r}^{\text{rep}}, \mathbf{\Phi}) \geq D(\mathbf{r}^{\text{obs}}, \mathbf{\Phi}) | \mathbf{r}^{\text{obs}}\Big) \\ &\approx \frac{1}{M} \sum_{i=1}^{M} \mathbb{1}_{\{D(r_i^{\text{rep}}, \mathbf{\Phi}^{(i)}) \geq D(\mathbf{r}^{\text{obs}}, \mathbf{\Phi}^{(i)})\}}. \end{aligned}$$

To interpret this quantity, note that if typical simulations are close to the observed data then we would expect the ppp-value to be around 0.5, while values close to 0 or 1 would indicate a poor model fit. We carried out this procedure with $M_1 = M = 100$ and obtained a value of 0.42, which is suggestive of a good model fit. A more accurate value could in principle be obtained using larger values of $M$ and $M_1$, but the procedure is highly time-consuming in practice due to the fact that we require simulated epidemics of a given final size.

As a final assessment of model fit, from 5000 simulations we found that on average 91% of those infected were FTC members, compared with 94% in the data, although we also

found that around 20% of those infected were from outside the compounds, compared to no such individuals in the data.

# 4 Discussion

*Transmission* The posterior estimates of the basic model parameters indicate clearly that the dominant mode of transmission was within-compound between individuals of the same confession. Our estimates of who-infected-whom align with this; for instance, from Figure 2, the vast majority of transmission events occurred within a compound in the most-likely infection pathway. The epidemiological investigation reported in Thompson and Foege (1968) found that spread within compounds, and within families in particular, appeared to drive the epidemic and that membership of FTC itself was not the primary transmission route. This is in agreement with our findings. As in Eichner and Dietz (2003), we found that infectiousness is markedly less during the fever period than the rash period, although our mean estimate of the reduction parameter $b$ is larger than their maximum likelihood value, which is most likely due to the skewed shape of the marginal posterior density.

*Reproduction numbers* Our posterior mean estimate of $R_0$ is close to 8. This is slightly larger than the Eichner and Dietz estimate (6.87) but underlines the potentially devastating nature of smallpox. Such values are radically different from those obtained using simpler models: for instance, assuming an SIR model in a homogeneously-mixing population of 120 FTC individuals typically results in $R_0$ estimates slightly larger than one, even allowing for the infection rate to vary with time (see e.g. O'Neill and Roberts (1999), Xu *et al.* (2016)). This highlights the importance of models which properly take population structure into account. As previously stated, $R_0$ can be interpreted as the average number of secondary cases produced by a single infective individual in a large susceptible population. For the Abakaliki data, such an interpretation is hard to apply directly since the compounds, wherein most transmission occurs, are small enough to provide a rapid saturation effect via the depletion of available susceptible individuals.

*Model adequacy* Our model appears to fit the data reasonably well, with the possible caveat that the model invariably predicts cases occurring outside of the compounds. The fact that the entire population of Abakaliki is rather unrealistically modelled as a homogeneously mixing population goes some way to explaining this; in particular, the potential for contacts between those inside and outside the compounds, and especially between FTC members and those outside the compounds, could well have been rather less than that assumed in the model. According to Thompson and Foege (1968), the FTC community was largely isolated from the community at large, although several of its adult members were involved in trading activities in and around Abakaliki. Consequently, a model in which some fraction of FTC members had contact with the outside community might be more realistic, although there are no data to directly inform this. Another aspect that is missing from our model is that of

age categories; Thompson and Foege (1968) state that the highest attack rates were among children. However, there do not appear to be sufficient data on compound composition to accurately incorporate age categories, and it seems likely that a model with age-specific transmission rates may be over-parameterised.

*Control measures and the end of the outbreak* It seems likely that the advent of control measures at time $t_q$ played a crucial role in bringing the outbreak to its conclusion rapdily. Under the model assumptions, control measures reduce the rash period from an average of 16 days to just 2 days, which in turn reduces the number of new infections. Interestingly, the posterior mean of $R_0$ after $t_q$ (i.e. $R_0$ with $\mu_R = \mu_Q = 2.0$) is around 1.5, but this in itself is insufficient to permit further large-scale spread due to the depletion of susceptibles within the compounds, and the fact that the epidemic in the population outside the compounds is sub-critical (i.e. the basic reproduction number is less than 1). Expanding the latter point, if we define pre- and post-control measure reproduction numbers for spread within compounds, FTC and the wider population (e.g. $R_a = (\mu_R + b\mu_F)\lambda_a$, etc.), then posterior mean estimates show that (i) within compounds, the epidemic is super-critical before and after $t_q$; (ii) with the FTC community, the epidemic switches from super- to sub-critical; (iii) in the wider population, the epidemic is always sub-critical. Despite this, increasing the value of $t_q$ in simulations was found to increase the outbreak size; for instance, setting $t_q$ to be 50, 100 and 200 gave mean outbreak sizes of around 24, 44 and 64, respectively. However, with no restrictions, we found the average outbreak size to be around 86, which underlines the fact that the epidemic was sub-critical in the wider population.

*Accuracy of the Eichner and Dietz likelihood approximation* It is of interest to see that our results are fairly similar to those obtained by Eichner and Dietz (2003). The most plausible explanation for this is the fact that distributions used for the length of time in each disease stage do not have particularly large variances, which in turn means that the model is not all that different to one in which all event times are assumed known. For such a model, the approximation method used by Eichner and Dietz gives the true likelihood, essentially because the distributions used to approximate uncertain event times collapse to point masses around the true values. A further point of interest is that the Eichner and Dietz approximation produces a likelihood function which is numerically but not analytically tractable, specifically because it involves integrals that must be evaluated numerically. Although this is sufficient for optimization purposes such as maximum likelihood, in practice such likelihood functions can be computationally prohibitive for use within MCMC algorithms since they must be repeatedly evaluated. It would therefore be of interest to develop analytically tractable approximate likelihood functions.

## 5 Acknowledgements

It is a pleasure to thank Martin Eichner for helpful discussions about this work. J.E.S. was funded by a UK Engineering and Physical Sciences Research Council studentship.

## References

Bailey, N. T. J. and Thomas, A. S. (1971). The estimation of parameters from population data on the general stochastic epidemic. *Theoretical Population Biology*, 2:253–270

Becker, N. (1976). Estimation for an epidemic model. *Biometrics*, 32:769–777

Boys, R. J. and Giles, P. R. (2007). Bayesian inference for stochastic epidemic models with time-homogeneous removal rates. *Journal of Mathematical Biology*, 55:223–247

Clancy, D. and O'Neill P. D. (2008). Bayesian estimation of the basic reproduction number in stochastic epidemic models. *Bayesian Analysis*, 3:737–757

Eichner, M. and Dietz, K. (2003). Transmission Potential of Smallpox: Estimates based on detailed data from an outbreak. *American Journal of Epidemiology*, 158(2):110–117

Eto, A., Saito, T., Yokote, H., Kurane, I. and Kanatani, Y. (2015) Recent advances in the study of live attenuated cell-cultured smallpox vaccine LC16m8. *Vaccine*, 33(45):6106–6111.

Gani, R. and Leach S. (2001). Transmission potential of smallpox in contemporary populations. *Nature*, 414:748–751

Gelman, A., Meng, X. L. and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica sinica*, 6(4):733–760

Gibson, G. J. and Renshaw E. (1998). Estimating parameters in stochastic compartmental models using Markov chain methods. *IMA Journal of Mathematics Applied in Medicine and Biology*, 15:19–40

Golightly A., Henderson, D. A. and Sherlock C. (2014). Delayed acceptance particle MCMC for exact inference in stochastic kinetic models. *Statistics and Computing*, 1–17

Halloran, M. E., Longini I. M., Nizam, A. and Yang, Y. (2002). Containing bioterrorist smallpox. *Science*, 298:1428–1432

Henderson, D. A. and Isao, A. (2014) The smallpox threat: A time to reconsider global policy. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, 12(3):117–121.

Huggins, R. M., Yip, P. S. F. and Lau, E. H. Y. (2004). A note on the estimation of the initial number of susceptible individuals in the general epidemic model. *Statistics and Probability Letters*, 67:321–330

Kypraios, T. (2009). A note on maximum likelihood estimation of the initial number of susceptibles in the general stochastic epidemic model. *Statistics and Probability Letters*, 79(18):1972–1976

Lau, E. H. Y. and Yip, S. F. (2008). Estimating the basic reproductive number in the general epidemic model with an unknown initial number of susceptible individuals. *Scandinavian Journal of Statistics*, 35:650–663

McKinley, T. J., Ross, J. V., Deardon, R. and Cook, A. R. (2014). Simulation-based Bayesian inference for epidemic models. *Computational Statistics and Data Analysis*, 71:434–447

Meltzer, M. E., Damon, I., LeDuc, J. M. and Millar, J. D. (2001). Modeling potential responses to smallpox as a bioterrorist weapon. *Emerging Infectious Diseases*, 7:959–969

Oh, C. (2014). Improved estimation of the initial number of susceptible individuals in the general stochastic epidemic model using penalized likelihood. *The Scientific World Journal*, Article ID 241687, 6 pages. doi:10.1155/2014/241687

O'Neill, P. D. and N. G. Becker (2001). Inference for an epidemic when susceptibility varies. *Biostatistics*, 2:99–108

O'Neill, P. D. and Roberts, G. O. (1999). Bayesian inference for partially observed stochastic epidemics. *Journal of the Royal Statistical Society, Series A*, 162:121–129

Ray, J. and Marzouk, Y. M. (2008). A Bayesian method for inferring transmission chains in a partially observed epidemic. *Proceedings of the Joint Statistical Meetings*, 3–7

Shanmugan, R. (2011). Spinned Poisson distribution with health management application. *Health Care Management Science*, 14:299–306

Thompson, D. and Foege, W. (1968). Faith Tabernacle smallpox epidemic. Abakaliki, Nigeria. *World Health Organization*, 3:1–9

World Health Organization (2015). The Independent Advisory Group on Public Health Implications of Synthetic Biology Technology Related to Smallpox. *World Health Organization*, WHO/HSE/PED/2015. 1.

Xiang, F. and Neal, P. (2014). Efficient MCMC for temporal epidemics via parameter reduction. *Computational Statistics and Data Analysis*, 80:240–250

Xu, X., Kypraios, T. and O'Neill, P. D. (2016) Bayesian non-parametric inference for stochastic epidemic models using Gaussian Processes *Biostatistics*, in press. doi: 10.1093/biostatistics/kxw011

Yip, P. S. F. (1989). Estimating the initial relative infection rate for a stochastic epidemic model. *Theoretical Population Biology*, 36:202–213